

Excess risk bounds for multitask learning with trace norm regularization

Andreas Maurer
Adalbertstr. 55, D-80799 München, Germany
am@andreas-maurer.eu

Massimiliano Pontil
Department of Computer Science
University College London
Malet Place London WC1E, UK
m.pontil@cs.ucl.ac.uk

January 15, 2013

Abstract

Trace norm regularization is a popular method of multitask learning. We give excess risk bounds with explicit dependence on the number of tasks, the number of examples per task and properties of the data distribution. The bounds are independent of the dimension of the input space, which may be infinite as in the case of reproducing kernel Hilbert spaces. A byproduct of the proof are bounds on the expected norm of sums of random positive semidefinite matrices with subexponential moments.

1 Introduction

A fundamental limitation of supervised learning is the cost incurred by the preparation of the large training samples required for good generalization. A potential remedy is offered by multi-task learning: in many cases, while individual sample sizes are rather small, there are samples to represent a large number of learning tasks, which share some constraining or generative property. This common property can be estimated using the entire collection of training samples, and if this property is sufficiently simple it should allow better estimation of the individual tasks from small individual samples.

The machine learning community has tried multi-task learning for many years (see [3, 4, 12, 13, 14, 20, 21, 26], contributions and references therein), but there are few theoretical investigations which clearly expose the conditions under which multi-task learning is preferable to independent learning. Following the seminal work of Baxter ([7, 8]) several authors have given generalization and

performance bounds under different assumptions of task-relatedness. In this paper we consider multi-task learning with trace-norm regularization (TNML), a technique for which efficient algorithms exist and which has been successfully applied many times (see e.g. [2, 4, 14, 15]).

In the learning framework considered here the inputs live in a separable Hilbert space H , which may be finite or infinite dimensional, and the outputs are real numbers. For each of T tasks an unknown input-output relationship is modeled by a distribution μ_t on $H \times \mathbb{R}$, with $\mu_t(X, Y)$ being interpreted as the probability of observing the input-output pair (X, Y) . We assume bounded inputs, for simplicity $\|X\| \leq 1$, where we use $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$ to denote euclidean norm and inner product in H respectively.

A predictor is specified by a weight vector $w \in H$ which predicts the output $\langle w, x \rangle$ for an observed input $x \in H$. If the observed output is y a loss $\ell(\langle w, x \rangle, y)$ is incurred, where ℓ is a fixed loss function on \mathbb{R}^2 , assumed to have values in $[0, 1]$, with $\ell(\cdot, y)$ being Lipschitz with constant L for each $y \in \mathbb{R}$. The expected loss or risk of weight vector w in the context of task t is thus

$$R_t(w) = \mathbb{E}_{(X,Y) \sim \mu_t} [\ell(\langle w, X \rangle, Y)].$$

The choice of a weight vector w_t for each task t is equivalent to the choice of a linear map $W : H \rightarrow \mathbb{R}^T$, with $(Wx)_t = \langle x, w_t \rangle$. We seek to choose W so as to (nearly) minimize the total average risk $R(W)$ defined by

$$R(W) = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{(X,Y) \sim \mu_t} [\ell(\langle w_t, X \rangle, Y)].$$

Since the μ_t are unknown, the minimization is based on a finite sample of observations, which for each task t is modelled by a vector \mathbf{Z}^t of n independent random variables $\mathbf{Z}^t = (Z_1^t, \dots, Z_n^t)$, where each $Z_i^t = (X_i^t, Y_i^t)$ is distributed according to μ_t . For most of this paper we make the simplifying assumption that all the samples have the same size n . With an appropriate modification of the algorithm defined below this assumption can be removed (see Remark 7 below). In a similar way the definition of $R(W)$ can be replaced by a weighted average which attribute greater weight to tasks which are considered more important. The entire multi-sample $(\mathbf{Z}^1, \dots, \mathbf{Z}^T)$ is denoted by $\bar{\mathbf{Z}}$.

A classical and intuitive learning strategy is empirical risk minimization. One decides on a constraint set $\mathcal{W} \subseteq \mathcal{L}(H, \mathbb{R}^T)$ for candidate maps and solves the problem

$$\hat{W}(\bar{\mathbf{Z}}) = \arg \min_{W \in \mathcal{W}} \hat{R}(W, \bar{\mathbf{Z}}),$$

where the average empirical risk $\hat{R}(W, \bar{\mathbf{Z}})$ is defined as

$$\hat{R}(W, \bar{\mathbf{Z}}) = \frac{1}{T} \sum_{t=1}^T \frac{1}{n} \sum_{i=1}^n \ell(\langle w_t, X_i^t \rangle, Y_i^t).$$

If the candidate set \mathcal{W} has the form $\mathcal{W} = \{x \mapsto Wx : (Wx)_t = \langle x, w_t \rangle, w_t \in \mathcal{B}\}$ where $\mathcal{B} \subseteq H$ is some candidate set of vectors, then this is equivalent to single task learning, solving for each task the problem

$$w_t(\mathbf{Z}_t) = \arg \min_{w \in \mathcal{B}} \frac{1}{n} \sum_{i=1}^n \ell(\langle w, X_i^t \rangle, Y_i^t).$$

For proper multi-task learning the set \mathcal{W} is chosen such that for a map W membership in \mathcal{W} implies some mutual dependence between the vectors w_t .

A good candidate set \mathcal{W} must fulfill two requirements: it must be large enough to contain maps with low risk and small enough that we can find such maps from a finite number of examples. The first requirement means that the risk of the best map W^* in the set,

$$W^* = \arg \min_{W \in \mathcal{W}} R(W),$$

is small. This depends on the set of tasks at hand and is largely a matter of domain knowledge. The second requirement is that the risk of the operator which we find by empirical risk minimization, $\hat{W}(\bar{\mathbf{Z}})$, is not too different from the risk of W^* , so that the excess risk

$$R(\hat{W}(\bar{\mathbf{Z}})) - R(W^*)$$

is small. Bounds on this quantity are the subject of this paper, and, as $R(\hat{W}(\bar{\mathbf{Z}}))$ is a random variable, they can only be expected to hold with a certain probability.

For multitask learning with trace-norm regularization (TNML) we suppose that \mathcal{W} is defined in terms of the trace-norm

$$\mathcal{W} = \left\{ W \in \mathbb{R}^{dT} : \|W\|_1 \leq B\sqrt{T} \right\}, \quad (1)$$

where $\|W\|_1 = \text{tr}((W^*W)^{1/2})$ and $B > 0$ is a regularization constant. The factor \sqrt{T} is an important normalization which we explain below. We will prove

Theorem 1 (i) For $\delta > 0$ with probability at least $1 - \delta$ in $\bar{\mathbf{Z}}$

$$R(\hat{W}) - R(W^*) \leq 2LB \left(\sqrt{\frac{\|C\|_\infty}{n}} + 5\sqrt{\frac{\ln(nT) + 1}{nT}} \right) + \sqrt{\frac{2\ln(2/\delta)}{nT}},$$

where $\|\cdot\|_\infty$ is the operator, or spectral norm, and C is the task averaged, uncentered data covariance operator

$$\langle Cv, w \rangle = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{(X,Y) \sim \mu_t} \langle v, X \rangle \langle X, w \rangle, \text{ for } w, v \in H.$$

(ii) Also with probability $1 - \delta$ in $\bar{\mathbf{Z}}$

$$R(\hat{W}) - R(W^*) \leq 2LB \left(\sqrt{\frac{\|\hat{C}\|_\infty}{n}} + \sqrt{\frac{2(\ln(nT) + 1)}{nT}} \right) + \sqrt{\frac{8 \ln(3/\delta)}{nT}},$$

with \hat{C} being the task averaged, uncentered empirical covariance operator

$$\langle \hat{C}v, w \rangle = \frac{1}{nT} \sum_{t=1}^T \sum_{i=1}^n \langle v, X_i^t \rangle \langle X_i^t, w \rangle, \text{ for } w, v \in H.$$

Remarks:

1. The first bound is distribution dependent, the second data-dependent.
2. Suppose that for an operator W all T column vectors w_t are equal to a common vector w , as might be the case if all the tasks T are equivalent. In this case increasing the number of tasks should not increase the regularizer. Since then $\|W\|_1 = \sqrt{T} \|w\|$ we have chosen the factor \sqrt{T} in (1). It allows us to consider the limit $T \rightarrow \infty$ for a fixed value of B .
3. In the limit $T \rightarrow \infty$ the bounds become

$$2LB \sqrt{\frac{\|C\|_\infty}{n}} \text{ or } 2LB \sqrt{\frac{\|\hat{C}\|_\infty}{n}} \text{ respectively.}$$

The limit is finite and it is approached at a rate of $\sqrt{\ln(T)/T}$.

4. If the mixture of data distributions is supported on a one dimensional subspace then $\|C\|_\infty = \mathbb{E} \|X\|^2$ and the bound is always worse than standard bounds for single task learning as in [6]. The situation is similar if the distribution is supported on a very low dimensional subspace. Thus, if learning is already easy, TNML will bring no benefit.
5. If the mixture of data distributions is uniform on an M -dimensional unit sphere in H then $\|C\|_\infty = 1/M$ and the corresponding term in the bound becomes small. Suppose now that for $W = (w_1, \dots, w_T)$ the w_t all are constrained to be unit vectors lying in some K -dimensional subspace of H , as might be the solution returned by a method of subspace learning [3]. If we choose $B = K^{1/2}$ then $W \in \mathcal{W}$, and our bound also applies. This subspace corresponds to the property shared among the tasks. The cost of its estimation vanishes in the limit $T \rightarrow \infty$ and the bound becomes

$$2L \sqrt{\frac{K}{nM}}.$$

K is proportional to the number of bits needed to communicate the utilized component of an input vector, given knowledge of the common subspace.

M is proportional to the number of bits to communicate an entire input vector. In this sense the quantity K/M can be interpreted as the ratio of the utilized information K to the available information M , as in [22]. If T and M are large and K is small the excess risk can be very small even for small sample sizes m . Thus, if learning is difficult (due to data of intrinsically high dimension) and the approximation error is small, then TNML is superior to single task learning.

6. An important example of the infinite dimensional case is given when H is the reproducing kernel Hilbert space H_κ generated by a positive semidefinite kernel $\kappa : Z \times Z \rightarrow \mathbb{R}$ where Z is a set of inputs. This setting is important because it allows to learn large classes of nonlinear functions. By the representer theorem for matrix regularizers [5] empirical risk minimization within the hypothesis space \mathcal{W} reduces to a finite dimensional problem in nT^2 variables.
7. The assumption of equal sample sizes for all tasks is often violated in practice. Let n_t be the number of examples available for the t -th task. The resulting imbalance can be compensated by a modification of the regularizer, replacing $\|W\|_1$ by a weighted trace norm $\|SW\|_1$, where the diagonal matrix $S = (s_1, \dots, s_T)$ weights the t -th task with

$$s_t = \sqrt{\frac{1}{n_t T} \sum_r n_r}.$$

where n_t is the size of the sample available for the t -th task. With this modification the Theorem holds with the average sample size $\bar{n} = (1/T) \sum n_t$ in place of n . In Section 5 we will prove this result, which then reduces to Theorem 1 when all the sample sizes are equal.

The proof of Theorem 1 is based on the well established method of Rademacher averages [6] and more recent advances on tail bounds for sums of random matrices, drawing heavily on the work of Ahlswede and Winter [1], Oliveira [24] and Tropp [27]. In this context two auxiliary results are established (Theorem 7 and Theorem 8 below), which may be of independent interest.

2 Earlier work.

The foundations to a theoretical understanding of multi-task learning were laid by J. Baxter in [8], where covering numbers are used to expose the potential benefits of multi-task and transfer learning. In [3] Rademacher averages are used to give excess risk bounds for a method of multi-task subspace learning. Similar results are obtained in [21]. [9] uses a special assumption of task-relatedness to give interesting bounds not on the average, but the maximal risk over the tasks.

A lot of important work on trace norm regularization concerns matrix completion, where a matrix is only partially observed and approximated (or under

certain assumptions even reconstructed) by a matrix of small trace norm (see e.g. [11], [25] and references therein). For $H = \mathbb{R}^d$ and $T \times d$ -matrices, this is somewhat related to the situation considered here, if we identify the tasks with the columns of the matrix in question, the input marginal as the uniform distribution supported on the basis vectors of \mathbb{R}^d and the outputs as defined by the matrix values themselves, without or with the addition of noise. One essential difference is that matrix completion deals with a known and particularly simple input distribution, which makes it unclear how bounds for matrix completion can be converted to bounds for multitask learning. On the other hand our bounds cannot be directly applied to matrix completion, because they assume a fixed number of revealed entries for each column.

Multitask learning is considered in [20], where special assumptions (coordinate-sparsity of the solution, restricted eigenvalues) are used to derive fast rates and the recovery of shared features. Such assumptions are absent in this paper, and [20] also considers a different regularizer.

[22] and [18] seem to be most closely related to the present work. In [22] the general form of the bound is very similar to Theorem 1. The result is dimension independent, but it falls short of giving the rate of $\sqrt{\ln(T)}/T$ in the number of tasks. Instead it gives $T^{-1/4}$.

[18] introduces a general and elegant method to derive bounds for learning techniques which employ matrix norms as regularizers. For $H = \mathbb{R}^d$ and applied to multi task learning and the trace-norm a data-dependent bound is given whose dominant term reads as (omitting constants and observing $\|W\|_1 \leq B\sqrt{T}$)

$$LB\sqrt{\max_i \|\hat{C}_i\|_\infty \frac{\ln \min\{T, d\}}{n}}, \quad (2)$$

where the matrix \hat{C}_i is the empirical covariance of the data for all tasks observed in the i -th observation

$$\hat{C}_i v = \frac{1}{T} \sum_t \langle v, X_i^t \rangle X_i^t.$$

The bound (2) does not paint a clear picture of the role of the number of tasks T . Using Theorem 8 below we can estimate its expectation and convert it into the distribution dependent bound with dominant term

$$LB\sqrt{\ln \min\{T, d\}} \left(\sqrt{\frac{\|C\|_\infty}{n}} + \sqrt{\frac{6 \ln(24nT^2) + 1}{nT}} \right). \quad (3)$$

This is quite similar to Theorem 1 (i). Because (2) is hinged on the i -th observation it is unclear how it can be modified for unequal sample sizes for different tasks. The principal disadvantage of (2) however is that it diverges in the simultaneous limit $d, T \rightarrow \infty$.

3 Notation and Tools

The letters H, H', H'' will denote finite or infinite dimensional separable real Hilbert spaces.

For a linear map $A : H \rightarrow H'$ we denote the adjoint with A^* , the range by $Ran(A)$ and the null space by $Ker(A)$. A is called compact if the image of the open unit ball of H under A is pre-compact (totally bounded) in H' . If $Ran(H)$ is finite dimensional then A is compact, finite linear combinations of compact linear maps and products with bounded linear maps are compact. A linear map $A : H \rightarrow H$ is called an operator and self-adjoint if $A^* = A$ and nonnegative (or positive) if it is self-adjoint and $\langle Ax, x \rangle \geq 0$ (or $\langle Ax, x \rangle > 0$) for all $x \in H, x \neq 0$, in which case we write $A \succeq 0$ (or $A \succ 0$). We use " \preceq " to denote the order induced by the cone of nonnegative operators.

For linear $A : H \rightarrow H'$ and $B : H' \rightarrow H''$ the product $BA : H \rightarrow H''$ is defined by $(BA)x = B(Ax)$. Then $A^*A : H \rightarrow H$ is always a nonnegative operator. We use $\|A\|_\infty$ for the norm $\|A\|_\infty = \sup \{\|Ax\| : \|x\| \leq 1\}$. We generally assume $\|A\|_\infty < \infty$.

If A is a compact and self-adjoint operator then there exists an orthonormal basis e_i of H and real numbers λ_i satisfying $|\lambda_i| \rightarrow 0$ such that

$$A = \sum_i \lambda_i Q_{e_i},$$

where Q_{e_i} is the operator defined by $Q_{e_i}x = \langle x, e_i \rangle e_i$. The e_i are eigenvectors and the λ_i eigenvalues of A . If f is a real function defined on a set containing all the λ_i a self-adjoint operator $f(A)$ is defined by

$$f(A) = \sum_i f(\lambda_i) Q_{e_i}.$$

$f(A)$ has the same eigenvectors as A and eigenvalues $f(\lambda_i)$. In the sequel self-adjoint operators are assumed to be either compact or of the form $f(A)$ with A compact (we will encounter no others), so that there always exists a basis of eigenvectors. A self-adjoint operator is nonnegative (positive) if all its eigenvalues are nonnegative (positive). If A is positive then $\ln(A)$ exists and has the property $\ln(A) \preceq \ln(B)$ whenever B is positive and $A \preceq B$. This property of operator monotonicity will be tacitly used in the sequel.

We write $\lambda_{\max}(A)$ for the largest eigenvalue (if it exists), and for nonnegative operators $\lambda_{\max}(\cdot)$ always exists and coincides with the norm $\|\cdot\|_\infty$.

A linear subspace $M \subseteq H$ is called invariant under A if $AM \subseteq M$. For a linear subspace $M \subseteq H$ we use M^\perp to denote the orthogonal complement $M^\perp = \{x \in H : \langle x, y \rangle = 0, \forall y \in M\}$. For a selfadjoint operator $Ran(A)^\perp = Ker(A)$. For a self-adjoint operator A on H and an invariant subspace M of A the trace $tr_M A$ of A relative to M is defined

$$tr_M A = \sum_i \langle Ae_i, e_i \rangle,$$

where $\{e_i\}$ is a orthonormal basis of M . The choice of basis does not affect the value of tr_M . For $M = H$ we just write tr without subscript. The trace-norm of any linear map from H to any Hilbert space is defined as

$$\|A\|_1 = \text{tr} \left((A^*A)^{1/2} \right).$$

If $\|A\|_1 < \infty$ then A is compact. If A is an operator and $A \succeq 0$ then $\|A\|_1$ is simply the sum of eigenvalues of A . In the sequel we will use Hoelder's inequality [10] for linear maps in the following form.

Theorem 2 *Let A and B be two linear maps $H \rightarrow \mathbb{R}^T$. Then $|\text{tr}(A^*B)| \leq \|A\|_1 \|B\|_\infty$.*

Rank-1 operators and covariance operators. For $w \in H$ we define an operator Q_w by

$$Q_w v = \langle v, w \rangle w, \text{ for } v \in H.$$

In matrix notation this would be the matrix ww^* . It can also be written as the tensor product $w \otimes w$. We apologize for the unusual notation Q_w , but it will save space in many of the formulas below. The covariance operators in Theorem 1 are then given by

$$C = \frac{1}{T} \sum_t \mathbb{E}_{(X,Y) \sim \mu_t} Q_X \text{ and } \hat{C} = \frac{1}{nT} \sum_{t,i} Q_{X_i^t}.$$

Here and in the sequel the Rademacher variables σ_i^t (or sometimes σ_i) are uniformly distributed on $\{0, 1\}$, mutually independent and independent of all other random variables, and \mathbb{E}_σ is the expectation conditional on all other random variables present. We conclude this section with two lemmata. Two numbers $p, q > 1$ are called conjugate exponents if $1/p + 1/q = 1$.

Lemma 3 (i) *Let p, q be conjugate exponents and $s, a \geq 0$. Then $(\sqrt{s+pa} - \sqrt{a})^2 \geq s/q$. (ii) For $a, b > 0$*

$$\min_{q, p > 1 \text{ and } 1/q + 1/p = 1} \sqrt{pa + qb} = \sqrt{a} + \sqrt{b}.$$

(iii) *and for $a, b > 0$ we have $2\sqrt{ab} \leq (p-1)a + (q-1)b$.*

Proof. For conjugate exponents p and q we have $p-1 = p/q$ and $q-1 = q/p$. Therefore $pa + qb - (\sqrt{a} + \sqrt{b})^2 = (\sqrt{pa/q} - \sqrt{qb/p})^2 \geq 0$, which proves (iii) and gives $\sqrt{pa + qb} \geq \sqrt{a} + \sqrt{b}$. Take $s = qb$, subtract \sqrt{a} and square to get (i). Set $p = 1 + \sqrt{b/a}$ and $q = 1 + \sqrt{a/b}$ to get (ii). ■

Lemma 4 *Let $a, c > 0, b \geq 1$ and suppose the real random variable $X \geq 0$ satisfies $\Pr\{X > pa + s\} \leq b \exp(-s/(cq))$ for all $s \geq 0$ and all conjugate exponents p and q . Then*

$$\sqrt{\mathbb{E}X} \leq \sqrt{a} + \sqrt{c(\ln b + 1)}.$$

Proof. We use partial integration.

$$\begin{aligned}\mathbb{E}X &\leq pa + qc \ln b + \int_{qc \ln b}^{\infty} \Pr\{X > pa + s\} ds \\ &\leq pa + qc \ln b + b \int_{qc \ln b}^{\infty} e^{-s/(cq)} ds = pa + q(c \ln b + 1).\end{aligned}$$

Take the square root of both sides and use Lemma 3 (ii) to optimize in p and q to obtain the conclusion. ■

4 Sums of random operators

In this section we prove two concentration results for sums of nonnegative operators with finite dimensional ranges. The first (Theorem 7) assumes only a weak form of boundedness, but it is strongly dimension dependent. The second result (Theorem 8) is the opposite. We will use the following important result of Tropp (Lemma 3.4 in [27]), derived from Lieb's concavity theorem (see [10], Section IX.6):

Theorem 5 *Consider a finite sequence A_k of independent, random, self-adjoint operators and a finite dimensional subspace $M \subseteq H$ such that $A_k M \subseteq M$. Then for $\theta \in \mathbb{R}$*

$$\mathbb{E} \operatorname{tr}_M \exp \left(\theta \sum_k A_k \right) \leq \operatorname{tr}_M \exp \left(\sum_k \ln \mathbb{E} e^{\theta A_k} \right).$$

A corollary suited to our applications is the following

Theorem 6 *Let A_1, \dots, A_N be of independent, random, self-adjoint operators on H and let $M \subseteq H$ be a nontrivial, finite dimensional subspace such that $\operatorname{Ran}(A_k) \subseteq M$ a.s. for all k .*

(i) *If $A_k \succeq 0$ a.s then*

$$\mathbb{E} \exp \left(\left\| \sum_k A_k \right\| \right) \leq \dim(M) \exp \left(\lambda_{\max} \left(\sum_k \ln \mathbb{E} e^{A_k} \right) \right).$$

(ii) *If the A_k are symmetrically distributed then*

$$\mathbb{E} \exp \left(\left\| \sum_k A_k \right\| \right) \leq 2 \dim(M) \exp \left(\lambda_{\max} \left(\sum_k \ln \mathbb{E} e^{A_k} \right) \right).$$

Proof. Let $A = \sum_k A_k$. Observe that $M^\perp \subseteq \operatorname{Ker}(A) \cap (\cup_k \operatorname{Ker}(A_k))$, and that M is a nontrivial invariant subspace for A as well as for all the A_k .

(i) Assume $A_k \succeq 0$. Then also $A \succeq 0$. Since $M^\perp \subseteq \text{Ker}(A)$ there is $x_1 \in M$ with $\|x_1\| = 1$ and $Ax_1 = \|A\| x_1$ (this also holds if $A = 0$, since M is nontrivial). Thus $e^A x_1 = e^{\|A\|} x_1$. Extending x_1 to a basis $\{x_i\}$ of M we get

$$e^{\|A\|} = \langle e^A x_1, x_1 \rangle \leq \sum_i \langle e^A x_i, x_i \rangle = \text{tr}_M e^A.$$

Theorem 5 applied to the matrices which represent A_k restricted to the finite dimensional invariant subspace M then gives

$$\begin{aligned} \mathbb{E} e^{\|A\|} &\leq \mathbb{E} \text{tr}_M e^A \leq \text{tr}_M \exp \left(\sum_k \ln (\mathbb{E} e^{A_k}) \right) \\ &\leq \dim(M) \exp \left(\lambda_{\max} \left(\sum_k \ln (\mathbb{E} e^{A_k}) \right) \right), \end{aligned}$$

where the last inequality results from bounding tr_M by $\dim(M) \lambda_{\max}$ and $\lambda_{\max}(\exp(\cdot)) = \exp(\lambda_{\max}(\cdot))$.

(ii) Assume that A_k is symmetrically distributed. Then so is A . Since $M^\perp \subseteq \text{Ker}(A)$ there is $x_1 \in M$ with $\|x_1\| = 1$ and either $Ax_1 = \|A\| x_1$ or $-Ax_1 = \|A\| x_1$, so that either $e^A x_1 = e^{\|A\|} x_1$ or $e^{-A} x_1 = e^{\|A\|} x_1$. Extending to a basis again gives

$$e^{\|A\|} \leq \langle e^A x_1, x_1 \rangle + \langle e^{-A} x_1, x_1 \rangle \leq \text{tr}_M e^A + \text{tr}_M e^{-A}.$$

By symmetric distribution we have

$$\mathbb{E} e^{\|A\|} \leq \text{tr}_M (\mathbb{E} e^A + \mathbb{E} e^{-A}) \leq 2 \mathbb{E} \text{tr}_M e^A.$$

Then continue as in case (ii). ■

The following is our first technical tool.

Theorem 7 *Let $M \subseteq H$ be a subspace of dimension d and suppose that A_1, \dots, A_N are independent random operators satisfying $A_k \succeq 0$, $\text{Ran}(A_k) \subseteq M$ a.s. and*

$$\mathbb{E} A_k^m \preceq m! R^{m-1} \mathbb{E} A_k \tag{4}$$

for some $R \geq 0$, all $m \in \mathbb{N}$ and all $k \in \{1, \dots, N\}$. Then for $s \geq 0$ and conjugate exponents p and q

$$\Pr \left\{ \left\| \sum_k A_k \right\|_\infty > p \left\| \mathbb{E} \sum_k A_k \right\|_\infty + s \right\} \leq \dim(M) e^{-s/(qR)}.$$

Also

$$\sqrt{\mathbb{E} \left\| \sum_k A_k \right\|_\infty^2} \leq \sqrt{\left\| \mathbb{E} \sum_k A_k \right\|_\infty^2} + \sqrt{R (\ln \dim(M) + 1)}.$$

Proof. Let θ be any number satisfying $0 \leq \theta < \frac{1}{R}$. From (4) we get for any $k \in \{1, \dots, N\}$

$$\begin{aligned} \mathbb{E}e^{\theta A_k} &= I + \sum_{m=1}^{\infty} \frac{\theta^m}{m!} \mathbb{E}A_k^m \preceq I + \sum_{m=1}^{\infty} (\theta R)^m (R^{-1} \mathbb{E}A_k) \\ &= I + \frac{\theta}{1 - R\theta} \mathbb{E}A_k \preceq \exp\left(\frac{\theta}{1 - R\theta} \mathbb{E}A_k\right). \end{aligned}$$

Abbreviate $\mu = \|\mathbb{E} \sum_k A_k\|_{\infty}$ and let $r = s + p\mu$ and set

$$\theta = \frac{1}{R} \left(1 - \sqrt{\frac{\mu}{r}}\right),$$

so that $0 \leq \theta < 1/R$. Applying the above inequality and the operator monotonicity of the logarithm we get for all k that $\ln \mathbb{E} \exp(\theta A_k) \preceq \theta/(1 - R\theta) \mathbb{E}A_k$. Summing this relation over k and passing to the largest eigenvalue yields

$$\lambda_{\max} \left(\sum_k \ln \mathbb{E} e^{\theta A_k} \right) \leq \frac{\theta \mu}{1 - R\theta}$$

Now we combine Markov's inequality with Theorem 6 (i) and the last inequality to obtain

$$\begin{aligned} \Pr \left\{ \left\| \sum A_k \right\|_{\infty} \geq r \right\} &\leq e^{-\theta r} \mathbb{E} \exp \left(\theta \left\| \sum A_k \right\| \right) \\ &\leq \dim(M) e^{-\theta r} \exp \left(\lambda_{\max} \left(\sum_k \ln \mathbb{E} e^{\theta A_k} \right) \right) \\ &\leq \dim(M) \exp \left(-\theta r + \frac{\theta}{1 - R\theta} \mu \right) \\ &= \dim(M) \exp \left(\frac{-1}{R} (\sqrt{r} - \sqrt{\mu})^2 \right). \end{aligned}$$

By Lemma 3 (i) $(\sqrt{r} - \sqrt{\mu})^2 = (\sqrt{s + p\mu} - \sqrt{\mu})^2 \geq s/q$, so this proves the first conclusion. The second follows from the first and Lemma 4. ■

The next result and its proof are essentially due to Oliveira ([24], Lemma 1, but see also [23]). We give a slightly more general version which eliminates the assumption of identical distribution and has smaller constants.

Theorem 8 *Let A_1, \dots, A_N be independent random operators satisfying $0 \preceq A_k \preceq I$ and suppose that for some $d \in \mathbb{N}$*

$$\dim \text{Span}(Ran(A_1), \dots, Ran(A_N)) \leq d \tag{5}$$

almost surely. Then

(i)

$$\Pr \left\{ \left\| \sum_k (A_k - \mathbb{E}A_k) \right\|_\infty > s \right\} \leq 4d^2 \exp \left(\frac{-s^2}{9 \left\| \sum_k \mathbb{E}A_k \right\|_\infty + 6s} \right).$$

(ii)

$$\Pr \left\{ \left\| \sum_k A_k \right\|_\infty > p \left\| \mathbb{E} \sum_k A_k \right\|_\infty + s \right\} \leq 4d^2 e^{-s/(6q)}$$

(iii)

$$\sqrt{\mathbb{E} \left\| \sum_k A_k \right\|_\infty} \leq \sqrt{\left\| \mathbb{E} \sum_k A_k \right\|_\infty} + \sqrt{6(\ln(4d^2) + 1)}$$

In the previous theorem the subspace M was deterministic and had to contain the ranges of *all* possible random realizations of the A_k . By contrast the span appearing in (5) is the random subspace spanned by a single random realization of the A_k . If all the A_k have rank one, for example, we can take $d = N$ and apply the present theorem even if each $\mathbb{E}A_k$ has infinite rank. This allows to estimate the empirical covariance in terms of the true covariance for a bounded data distribution in an infinite dimensional space.

Proof. Let $0 \leq \theta < 1/4$ and abbreviate $A = \sum_k A_k$. A standard symmetrization argument (see [19], Lemma 6.3) shows that

$$\mathbb{E} e^{\theta \|A - \mathbb{E}A\|} \leq \mathbb{E} \mathbb{E}_\sigma \exp \left(2\theta \left\| \sum_k \sigma_k A_k \right\| \right),$$

where the σ_k are Rademacher variables and \mathbb{E}_σ is the expectation conditional on the A_1, \dots, A_N . For fixed A_1, \dots, A_N let M be the linear span of their ranges, which has dimension at most d and also contains the ranges of the symmetrically distributed operators $2\theta\sigma_k A_k$. Invoking Theorem 6 (ii) we get

$$\begin{aligned} \mathbb{E}_\sigma \exp \left(2\theta \left\| \sum_k \sigma_k A_k \right\| \right) &\leq 2d \exp \left(\lambda_{\max} \left(\sum_k \ln \mathbb{E}_\sigma e^{2\theta\sigma_k A_k} \right) \right) \\ &\leq 2d \exp \left(2\theta^2 \left\| \sum_k A_k^2 \right\| \right) \leq 2d \exp (2\theta^2 \|A\|). \end{aligned}$$

The second inequality comes from $\mathbb{E}_\sigma e^{2\theta\sigma_k A_k} = \cosh(2\theta A_k) \preceq e^{2\theta^2 A_k^2}$, and the fact that for positive operators λ_{\max} and the norm coincide. The last inequality follows from the implications $0 \preceq A_k \preceq I \implies A_k^2 \preceq A_k \implies \sum_k A_k^2 \preceq \sum_k A_k \implies \left\| \sum_k A_k^2 \right\| \leq \|A\|$. Now we take the expectation in A_1, \dots, A_N . Together with the previous inequalities we obtain

$$\begin{aligned} \mathbb{E} e^{\theta \|A - \mathbb{E}A\|} &\leq 2d \mathbb{E} e^{2\theta^2 \|A\|} \leq 2d \mathbb{E} e^{2\theta^2 \|A - \mathbb{E}A\|} e^{2\theta^2 \|\mathbb{E}A\|} \\ &\leq 2d \left(\mathbb{E} e^{\theta \|A - \mathbb{E}A\|} \right)^{2\theta} e^{2\theta^2 \|\mathbb{E}A\|}. \end{aligned}$$

The last inequality holds by Jensen's inequality since $\theta < 1/4 < 1/2$. Dividing by $(\mathbb{E} \exp(\theta \|A - \mathbb{E}A\|))^{2\theta}$, taking the power of $1/(1-2\theta)$ and multiplying with $e^{\theta s}$ gives

$$\Pr\{\|A - \mathbb{E}A\| > s\} \leq e^{-\theta s} \mathbb{E} e^{\theta \|A - \mathbb{E}A\|} \leq (2d)^{1/(1-2\theta)} \exp\left(\frac{2\theta^2}{1-2\theta} \|\mathbb{E}A\| - \theta s\right).$$

Since $\theta < 1/4$, we have $(2d)^{1/(1-2\theta)} < (2d)^2$. Substitution of $\theta = s/(6\|\mathbb{E}A\| + 4s) < 1/4$ together with some simplifications gives (i).

It follows from elementary algebra that for $\delta > 0$ with probability at least $1 - \delta$ we have

$$\begin{aligned} \|A\| &\leq \|\mathbb{E}A\| + 2\sqrt{\|\mathbb{E}A\|} \sqrt{\frac{9}{4} \ln(4d^2/\delta) + 6 \ln(4d^2/\delta)} \\ &\leq p \|\mathbb{E}A\| + 6q \ln(4d^2/\delta), \end{aligned}$$

where the last line follows from $(9/4) < 6$ and Lemma 3 (iii). Equating the second term in the last line to s and solving for the probability δ we obtain (ii), and (iii) follows from Lemma 4. ■

5 Proof of Theorem 1

We prove the excess risk bound for heterogeneous sample sizes with the weighted trace norm as in Remark 7 following the statement of Theorem 1. The sample size for the n -th task is thus n_t and we abbreviate \bar{n} for the average sample size, $\bar{n} = (1/T) \sum_t n_t$, so that $\bar{n}T$ is the total number of examples. The class of linear maps W considered is

$$\mathcal{W} = \left\{ W \in \mathbb{R}^{dT} : \|SW\|_1 \leq B\sqrt{T} \right\},$$

with $S = (s_1, \dots, s_T)$ and $s_t = \sqrt{\bar{n}/n_t}$. With \mathcal{W} so defined we will prove the inequalities in Theorem 1 with n replaced by \bar{n} . The result then reduces to Theorem 1 if all the sample sizes are equal.

The first steps in the proof follow a standard pattern. We write

$$\begin{aligned} &R(\hat{W}) - R(W^*) \\ &= \left[R(\hat{W}) - \hat{R}(\hat{W}, \bar{\mathbf{Z}}) \right] + \left[\hat{R}(\hat{W}, \bar{\mathbf{Z}}) - \hat{R}(W^*, \bar{\mathbf{Z}}) \right] + \left[\hat{R}(W^*, \bar{\mathbf{Z}}) - R(W^*) \right]. \end{aligned}$$

The second term is always negative by the definition of \hat{W} . The third term depends only on W^* . Using Hoeffding's inequality [16] it can be bounded with probability at least $1 - \delta$ by $\sqrt{\ln(1/\delta)/(2\bar{n}T)}$. There remains the first term which we bound by

$$\sup_{W \in \mathcal{W}} R(W) - \hat{R}(W).$$

It has by now become a standard technique (see [6]) to show that this quantity is with probability at least $1 - \delta$ bounded by

$$\mathbb{E}_{\bar{\mathbf{Z}}} \mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}}) + \sqrt{\frac{\ln(1/\delta)}{2\bar{n}T}} \quad (6)$$

or

$$\mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}}) + \sqrt{\frac{9 \ln(2/\delta)}{2\bar{n}T}}, \quad (7)$$

where the empirical Rademacher complexity $\mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}})$ is defined for a multi-sample $\bar{\mathbf{Z}}$ with values in $(H \times \mathbb{R})^{nT}$ by

$$\mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}}) = \frac{2}{T} \mathbb{E}_{\sigma} \sup_{W \in \mathcal{W}} \sum_{t=1}^T \frac{1}{n_t} \sum_{i=1}^{n_t} \sigma_i^t \ell(\langle w_t, X_i^t \rangle, Y_i^t).$$

Standard results on Rademacher averages allow us to eliminate the Lipschitz loss functions and give us

$$\begin{aligned} \mathcal{R}(\mathcal{W}, \bar{\mathbf{z}}) &\leq \frac{2L}{T} \mathbb{E}_{\sigma} \sup_{W \in \mathcal{W}} \sum_{t=1}^T \sum_{i=1}^{n_t} \sigma_i^t \langle w_t, X_i^t / n_t \rangle \\ &= \frac{2L}{T} \mathbb{E}_{\sigma} \sup_{W \in \mathcal{W}} \text{tr}(W^* D) = \frac{2L}{T} \mathbb{E}_{\sigma} \sup_{W \in \mathcal{W}} \text{tr}(W^* S S^{-1} D), \end{aligned}$$

where the random operator $D : H \rightarrow \mathbb{R}^T$ is defined for $v \in H$ by $(Dv)_t = \langle v, \sum_{i=1}^{n_t} \sigma_i^t X_i^t / n_t \rangle$, and the diagonal matrix S is as above. Hölder's and Jensen's inequalities give

$$\begin{aligned} \mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}}) &\leq \frac{2L}{T} \sup_{W \in \mathcal{W}} \|SW\|_1 \mathbb{E}_{\sigma} \|S^{-1}D\|_{\infty} = \frac{2LB}{n\sqrt{T}} \mathbb{E}_{\sigma} \|S^{-1}D\|_{\infty} \\ &\leq \frac{2LB}{\sqrt{T}} \sqrt{\mathbb{E}_{\sigma} \|D^* S^{-2} D\|_{\infty}}. \end{aligned}$$

Let V_t be the random vector $V_t = \sum_{i=1}^{n_t} \sigma_i^t X_i^t / (s_t n_t)$ and recall that the induced rank-one operator Q_{V_t} is defined by $Q_{V_t} v = \langle v, V_t \rangle V_t = (1/(\bar{n}n_t)) \sum_{ij} \langle v, \sigma_i^t X_i^t \rangle \sigma_j^t X_j^t$. Then $D^* S^{-2} D = \sum_{t=1}^T Q_{V_t}$, so we obtain

$$\mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}}) \leq \frac{2LB}{\sqrt{T}} \sqrt{\mathbb{E}_{\sigma} \left\| \sum_t Q_{V_t} \right\|_{\infty}}$$

as the central object which needs to be bounded.

Observe that the range of any Q_{V_t} lies in the subspace

$$M = \text{Span} \{X_i^t : 1 \leq t \leq T \text{ and } 1 \leq i \leq n_t\}$$

which has dimension $\dim M \leq \bar{n}T < \infty$. We can therefore pull the expectation inside the norm using Theorem 7 if we can verify a subexponential bound (4) on the moments of the Q_{V_t} . This is the content of the following lemma.

Lemma 9 Let x_1, \dots, x_n be in H and satisfy $\|x_i\| \leq b$. Define a random vector by $V = \sum_i \sigma_i x_i$. Then for $m \geq 1$

$$\mathbb{E}[(Q_V)^m] \leq m! (2nb^2)^{m-1} \mathbb{E}[Q_V].$$

Proof. Let $K_{m,n}$ be the set of all sequences (j_1, \dots, j_{2m}) with $j_k \in \{1, \dots, n\}$, such that each integer in $\{1, \dots, n\}$ occurs an even number of times. It is easily shown by induction that the number of sequences in $K_{m,n}$ is bounded by

$$|K_{m,n}| \leq (2m-1)!! n^m,$$

where $(2m-1)!! = \prod_{i=1}^m (2i-1) \leq m! 2^{m-1}$.

Now let $v \in H$ be arbitrary. By the definition of V and Q_V we have for any $v \in H$ that

$$\langle \mathbb{E}[(Q_V)^m] v, v \rangle = \sum_{j_1, \dots, j_{2m}=1}^n \mathbb{E}[\sigma_{j_1} \sigma_{j_2} \cdots \sigma_{j_{2m}}] \langle v, x_{j_1} \rangle \langle x_{j_2}, x_{j_3} \rangle \cdots \langle x_{j_{2m}}, v \rangle.$$

The properties of independent Rademacher variables imply that $\mathbb{E}[\sigma_{j_1} \sigma_{j_2} \cdots \sigma_{j_{2m}}] = 1$ if $j \in K_{m,n}$ and zero otherwise. For $m = 1$ this shows $\langle \mathbb{E}[(Q_V)^m] v, v \rangle = \langle \mathbb{E}[Q_V] v, v \rangle = \sum_j \langle v, x_j \rangle^2$. For $m > 1$, since $\|x_i\| \leq b$ and by two applications of the Cauchy-Schwarz inequality

$$\begin{aligned} \langle \mathbb{E}[(Q_V)^m] v, v \rangle &= \sum_{\mathbf{j} \in K_{m,n}} \langle v, x_{j_1} \rangle \langle x_{j_2}, x_{j_3} \rangle \cdots \langle x_{j_{2m}}, v \rangle \\ &\leq b^{2(m-1)} \sum_{\mathbf{j} \in K_{m,n}} |\langle v, x_{j_1} \rangle| |\langle x_{j_{2m}}, v \rangle| \\ &\leq b^{2(m-1)} \left(\sum_{\mathbf{j} \in K_{m,n}} \langle v, x_{j_1} \rangle^2 \right)^{1/2} \left(\sum_{\mathbf{j} \in K_{m,n}} \langle v, x_{j_{2m}} \rangle^2 \right)^{1/2} \\ &= b^{2(m-1)} \sum_j \langle v, x_j \rangle^2 \cdot \sum_{\mathbf{j} \in K_{m,n} \text{ such that } j_1=j} 1 \\ &= \langle \mathbb{E}[Q_V^m] v, v \rangle \times (2m-1)!! n^{m-1} b^{2(m-1)} \\ &\leq m! (2nb^2)^{m-1} \langle \mathbb{E}[Q_V] v, v \rangle. \end{aligned}$$

The conclusion follows since for self-adjoint matrices $(\forall v, \langle Av, v \rangle \leq \langle Bv, v \rangle) \implies A \preceq B$. ■

If we apply this lemma to the vectors V_t defined above with $b = 1/(s_t n_t)$, using $s_t^2 n_t = \bar{n}$, we obtain

$$\mathbb{E}[(Q_{V_t})^m] \leq m! \left(\frac{2}{s_t^2 n_t} \right)^{m-1} \mathbb{E}[Q_V] = m! \left(\frac{2}{\bar{n}} \right)^{m-1} \mathbb{E}[Q_V],$$

Applying the last conclusion of Theorem 7 with $R = 2/\bar{n}$ and $d = \bar{n}T$ now yields

$$\sqrt{\mathbb{E}_\sigma \left\| \sum Q_{V_t} \right\|_\infty} \leq \sqrt{\left\| \sum \mathbb{E}_\sigma Q_{V_t} \right\|_\infty} + \sqrt{\frac{2}{\bar{n}} (\ln(\bar{n}T) + 1)},$$

and since $\sum_t \mathbb{E}_\sigma Q_{V_t} = \sum_t (1/\bar{n}) \sum_i (1/n_t) Q_{X_i^t} = T\hat{C}/\bar{n}$ we get

$$\begin{aligned} \mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}}) &\leq \frac{2LB}{\sqrt{T}} \sqrt{\mathbb{E}_\sigma \left\| \sum_t Q_{V_t} \right\|_\infty} \\ &\leq 2LB \left(\sqrt{\frac{\|\hat{C}\|_\infty}{\bar{n}}} + \sqrt{\frac{2(\ln(\bar{n}T) + 1)}{\bar{n}T}} \right). \end{aligned} \quad (8)$$

Together with (7) and the initial remarks in this section this proves the second part of Theorem 1.

To obtain the first assertion we take the expectation of (8) and use Jensen's inequality, which then confronts us with the problem of bounding $\mathbb{E} \|\hat{C}\|_\infty$ in terms of $\|C\|_\infty = \|\mathbb{E}\hat{C}\|_\infty$. Note that $\bar{n}T\hat{C} = \sum_t \sum_{i=1}^{n_t} Q_{X_i^t}$. Here Theorem 7 doesn't help because the covariance may have infinite rank, so that we cannot find a finite dimensional subspace containing the ranges of all the $Q_{X_i^t}$. But since $\|X_i^t\| \leq 1$ all the $Q_{X_i^t}$ satisfy $0 \preceq Q_{X_i^t} \preceq I$ and are rank-one operators, we can invoke Theorem 8 with $d = \bar{n}T$. This gives

$$\sqrt{\mathbb{E} \|\hat{C}\|_\infty} \leq \sqrt{\|C\|_\infty} + \sqrt{\frac{6(\ln(4\bar{n}T) + 1)}{\bar{n}T}},$$

and from (8) and Jensen's inequality and some simplifications we obtain

$$\begin{aligned} \mathbb{E} \mathcal{R}(\mathcal{W}, \bar{\mathbf{Z}}) &\leq 2LB \left(\sqrt{\frac{\mathbb{E} \|\hat{C}\|_\infty}{\bar{n}}} + \sqrt{\frac{2(\ln(\bar{n}T) + 1)}{\bar{n}T}} \right) \\ &\leq 2LB \left(\sqrt{\frac{\|C\|_\infty}{\bar{n}}} + 5\sqrt{\frac{\ln(\bar{n}T) + 1}{\bar{n}T}} \right), \end{aligned}$$

which, together with (6), gives the first assertion of Theorem 1.

A similar application of Theorem 8 applied to the bound (2) in [18] yields the bound (3).

References

- [1] R. Ahlswede, A. Winter. Strong converse for identification via quantum channels, *IEEE Trans. Inf. Theory* 48(3), 569–579, 2002.

- [2] Y. Amit, M. Fink, N. Srebro, S. Ullman. Uncovering Shared Structures in Multiclass Classification. *24th International Conference on Machine Learning (ICML)*, 2007.
- [3] R. K. Ando, T. Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data, *Journal of Machine Learning Research*, 6: 1817–1853, 2005.
- [4] A. Argyriou, T. Evgeniou, M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3): 243–272, 2008
- [5] A. Argyriou, C.A. Micchelli, M. Pontil. When is there a representer theorem? Vector versus matrix regularizers. *Journal of Machine Learning Research*, 10:2507–2529, 2009.
- [6] P. L. Bartlett and S. Mendelson. Rademacher and Gaussian Complexities: Risk Bounds and Structural Results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [7] J. Baxter. Theoretical Models of Learning to Learn. In *Learning to Learn*, S.Thrun, L.Pratt Eds. Springer 1998
- [8] J. Baxter. A Model of Inductive Bias Learning, *Journal of Artificial Intelligence Research* 12:149–198, 2000.
- [9] S. Ben-David and R. Schuller. Exploiting task relatedness for multiple task learning. In *COLT 03*, 2003.
- [10] R. Bhatia. *Matrix Analysis*. Springer, 1997.
- [11] E. Candès and T. Tao. The power of convex relaxation: Near optimal matrix completion. *IEEE Trans. Inform. Theory*, 56(5):2053–2080, 2009.
- [12] R. Caruana. Multitask Learning. *Machine Learning*, 28(1):41–75, 1997.
- [13] G. Cavallanti, N. Cesa-Bianchi, and C. Gentile. Linear algorithms for online multitask classification. *Journal of Machine Learning Research*, 11:2597–2630, 2010.
- [14] T. Evgeniou, C. Micchelli and M. Pontil. Learning multiple tasks with kernel methods. *Journal of Machine Learning Research*, 6:615–637, 2005.
- [15] Z. Harchaoui, M. Douze, M. Paulin, M. Dudik, J. Malick. Large-scale classification with trace-norm regularization. *IEEE Conference on Computer Vision & Pattern Recognition (CVPR)*, 2012.
- [16] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
- [17] D. Hsu, S. Kakade, and T. Zhang. Dimension-free tail inequalities for sums of random matrices. *arXiv:1104.1672*, 2011.

- [18] S. M. Kakade, S. Shalev-Shwartz, A. Tewari. Regularization Techniques for Learning with Matrices. *Journal of Machine Learning Research* 13:1865–1890, 2012.
- [19] M. Ledoux, M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, Berlin, 1991.
- [20] K. Lounici, M. Pontil, A.B. Tsybakov and S. van de Geer. Oracle inequalities and optimal inference under group sparsity. *Annals of Statistics*. 2012.
- [21] A. Maurer. Bounds for linear multi-task learning. *Journal of Machine Learning Research*, 7:117–139, 2006.
- [22] A. Maurer. The Rademacher complexity of linear transformation classes. *Colt 2006*, Springer, 65–78, 2006.
- [23] S. Mendelson and A. Pajor. On singular values of matrices with independent rows. *Bernoulli* 12(5): 761–773, 2006.
- [24] R.I. Oliveira. Sums of random Hermitian matrices and an inequality by Rudelson, *Electron. Commun. Probab.* 15, 203–212, 2010.
- [25] O. Shamir and S. Shalev-Shwartz. Collaborative Filtering with the trace norm: Learning, bounding and transducing. *24th Annual Conference on Learning Theory (COLT)*, 2011.
- [26] S. Thrun and L. Pratt. *Learning to Learn*. Springer 1998.
- [27] J. Tropp. User-friendly tail bounds for sums of random matrices, *Foundations of Computational Mathematics*, 12:389–434, 2012.